

Distributed Accounting on the Grid

William Thigpen
NASA IPG

bthigpen@arc.nasa.gov

Thomas J. Hacker
University
of Michigan

hacker@umich.edu

Laura F. McGinnis
Pittsburgh
Supercomputer Center
lfm1@psc.edu

Brian D. Athey
University
of Michigan
bleu@umich.edu

The Internet has been engineered over the last thirty years to interconnect devices across the globe in an adaptable and fault-tolerant manner. Along with the development of the Internet, a suite of distributed applications ranging from electronic mail to the World Wide Web that rely upon the global Internet have grown in use and scope in parallel with the universal deployment and use of the Internet.

By the late 1990's, the Internet was adequately equipped to move vast amounts of data between HPC systems, and efforts were initiated to link together the national infrastructure of high performance computational and data storage resources together into a general computational utility "grid", analogous to the national electrical power grid infrastructure.

The purpose of the computational grid is to provide dependable, consistent, pervasive, and inexpensive access to computational resources for the computing community in the form of a computing utility [1].

This paper presents a fully distributed view of Grid usage accounting and a methodology for allocating Grid computational resources for use on a Grid computing system. The Grid will contain a large number of unconnected sites, and these sites will need to exchange accounting and bid/quote information.

Specifically, three issues are being addressed by the Global Grid Forum Accounting Working Group: mapping resource usage to GRID users; defining a usage economy or methods for resource exchange, and describing implementation standards that minimize and compartmentalize the tasks required for a site to participate in GRID accounting.

For an accounting system to be functional in a GRID environment, it needs to be decentralized, scalable and flexible. It must have a minimum impact on local accounting and should not make any limiting assumptions about whether accounting is done by user, group, project, or site. The requirements on the remote site will be to track the resources used by the requesting job and then pass this information back to the requesting site in some standardized format. At the requesting site, the information can then be accrued as needed for local requirements. A distributed allocation and accounting approach, using a consumer/supplier or client/server structure will work across multiple sites and satisfy the needs of the participating administrative and policy domains.

Mapping Usage to Users

The current situation at most potential GRID sites is that to run jobs on a machine, the user needs to have a local user account on that machine. Unfortunately, as the GRID grows in the number of sites and users, this method of establishing access to resources will not scale. For example, at the University of Michigan, over 120,000 users are registered and a significant amount of time and energy is spent managing this registry. As the GRID grows beyond this scale, continued reliance on the existence of a local user account would engender the need to create a centralized bureaucracy to manage this registry, which is antithetical to stated GRID goals.

It should be noted that if a site requires users to have local accounts for remote execution, then the site might not be able to use the full capabilities of the GRID. The GRID needs to be a fluid environment where sites can exchange cycles and provide access to users of other trusted participating GRID sites. The overhead and time delay in requiring local user accounts could easily become the critical bottleneck in this process.

Distributed accounting on the GRID assumes the existence of authentication and authorization mechanisms which securely and accurately establish the identity and credentials of user requesting access to GRID resources¹. Once identity and credentials have been established, distributed accounting methods must be able to map GRID resource usage to the requesting user. Since it has already been established that local user accounts are not feasible in a GRID environment, various methods of "accountless accounting" need to be investigated.

The Free Market Economy Model

In a free market economy, the allocation of resources is determined solely by supply and demand. Ideally, supply and demand are not subject to regulation other than normal competition, but property rights are allocated and upheld so that trade can occur.² A free market system provides the following benefits:

¹ Refer to the Security GRID working group for work being done on authorization and authentication.

² Based on definitions from The Penguin Dictionary of Economics, edited by Graham Bannock, R. E. Baxter, and Evan Davis; 1998 (<http://www.xrefer.com>).

Suppliers	Consumers
Resource Control: Each supplier site has control over the set of resources and the quantity of each resource that it chooses to make available on the GRID.	Resource Selection: Consumers can choose from a variety of resources that might not otherwise be available to them.
Price: Each supplier site can modify the set of resources and the rates for each resource as needed.	Value: The costs incurred in utilizing resources from various suppliers can be compared prior to submitting a request for resources.
Implementation: Each site can implement as complex or as simple an accounting system as needed. New accounting systems can be easily prototyped. Supplier sites can change the way they charge as they desire with minimum impact or requirements on other sites.	Implementation: Resource requests can be submitted independent of implementation details. The consumer only needs to know a standard method for requesting resources and compensating the resource supplier.
Autonomy: Supplier sites need not agree nor negotiate the relative value of their resources as a prerequisite to making those resources available on the GRID.	Independence: Consumers can compare resources available at various sites and select those that best meet their needs.
Exchange	
The free market model provides an automatic way to regulate the utilization of site resources by external resource consumers	

In the context of the GRID, certain fundamental concepts must be defined for resources to be equitably and efficiently allocated and utilized:

- **Supplier:** A provider of GRID resources
- **Consumer:** A user of GRID resources
- **Value:** A measurement of the usage of GRID resources. In the consumer's perspective, this could be seen as *cost* or *price*.
- **Exchange:** The act of utilizing GRID resources provided by a GRID supplier and received by a GRID consumer

Supplier Sites: The GRID resource supplier must be able to provide its resource rates, quotes for resource requests, and resource usage. The following mechanisms should be implemented at a GRID resource supplier site:

- **GRID Resource Provider Rates:** The GRID resource provider should have a way to set and maintain the rates for the use of its resources. There is no need for an agreement on how this information is stored or provided.
- **Provide Quotes for resource allocation requests:** An agreement on the format of this message must be a deliverable for the GRID Forum. The response to the resource allocation quote request will provide a cost for the requested resources. The final charge for the resource usage by the job should not exceed the quote if the job resource requirements did not exceed the estimates provided for the quote. The response to the resource quote request will contain the requester's authorization identifier, an expiration date/time that describe when the quote expires, and a server unique identifier. The resource utilization quote provided in response to a resource allocation request will be a total cost and will not be broken down by resource categories. If the request stipulated a range of charges, all ranges will be provided with a separate unique ID.
- **Track Resource Utilization:** Each GRID resource provider can choose to gather information on resources consumers by local and remote users. GRID resource providers must (it's in their own interest) collect information on GRID credits collected from resource consumers. This has to happen, but does not need to be specified in the GRID community. Each site must have the access and ability to track the information it will charge for against the particular job request.
- **Job Account Information:** The functionality required to package and transfer the resources utilized by a resource consumer must be defined and agreed upon by all GRID participants. For maximum flexibility, sites must be able to provide an accounting record (either pull or push). When the job completes

(normally or abnormally), the accounting information is gathered and sent back to the resource-consuming site. This accounting should be broken down by resource category and must include the requesting site's unique ID. This information must be delivered. If for some reason the delivery fails, the information must go to the supplier's accounting authority to handle manually.

Consumer Sites: The GRID resource consumer must be able to obtain quotes for future resource consumption and either request that the resource consuming job be executed or inform the resource provider that the resource quote was rejected.

➤ *Resource Usage Quote Query:* This should be a request in a common agreed upon format that specifies the resources requested. This resource quote request does not obligate the requester to use the requested services, it is simply a mechanism that the potential resource consumer can use to ascertain potential costs for utilizing the resources that the resource provider can provide. The resource quote request should have a requesting site unique identifier and a description of the resources required. The resource quote request can request a range of charges based upon additional qualifiers such as quality of service if provided by the resource provider site.

➤ *Accountable Resource Use Request:* If a resource consuming entity decides to use a resource provider site the resource consuming request will include a unique requester ID and will include the server ID associated with the resource quote provided to the resource consuming entity.

➤ *Resource Request Quote Cancellation:* Although not required, it is suggested if the requesting site decides to use the successful bidder for a job, a cancellation should be sent to all the resource providers that provided a quote whose resources are not going to be used. This would include canceling unused resource requests from the resource provider site that won the bid. If this cancellation were not sent, the reservation would be removed automatically when the quote expires.

Valuing Resources : The local resource provider determines the base value of resources within their administrative purview. This resource valuation can be used as a mechanism to attract or deter external users by utilizing the laws of supply and demand. Submitted jobs must therefore contain sufficient resource requirement information to allow local resource allocation software to determine the cost of the local resources that will be consumed by the job. The local authority will also need to decide, for their administrative purview, if a remote user is required to

have a local account to utilize local resources. If local resources are provided to remote users without local accounts/accounting, the local resource provider must provide a full accounting of each resource used and the costs charged for each resource for the job. This accounting can be performed immediately (e.g. at the completion of the job), later (i.e. when the accounting software is run), or upon request from the requesting site.

The rates determined by local resource providers for resources, while flexible, must be made available to a potential GRID user upon request for a quote. Resource quotes should contain a time frame for which the resource quote is valid. The quote process will facilitate an open bidding process for resources that will allow the user to comparison shop. This raises the additional question of how to release a quote that has not been accepted.

Functionality and Methodology

Chargeable Items. These are some of the major types of resources accounted for on the GRID. As a start, the following resource items are submitted for possible inclusion:

- Charge per CPU billing unit
- Charge by wall clock or usage billing unit
- Charge for memory
- If usage is tied to CPUs, amount / CPU
- Charge per megabyte of on-line storage
- Premium rate(s) for special handling
- Higher job queue priority within a job class
- Charge for network bandwidth usage (if bandwidth is pre-allocated and reserved)
- Special Application Charges
- Local consultant, programmer or administrator charges for time utilized beyond operation duties to provide assistance
- Transportable media charges (if data is moved to tape and shipped to requestor)

This list will grow, but the fact that an item is on the list does not mean that any resource provider must charge for it. For instance, a site may decide that it will only charge for CPU utilization. Conversely, this list is not exhaustive. Supplier sites who calculate "usage" using resource metrics not included here are welcome to define their "charges" to meet the unique requirements of their site or particular resource. Ultimately, the only requirement is that the resource usage be presented to the "consumer" in an understandable and decomposable fashion – the user needs to know what the measures are for using a site's resources so that an informed decision can be made before submitting a job.

Conflict Resolution: Each site must implement and publish its conflict resolution procedures for disputes over charges incurred. An overall procedure establishing minimum resolution standards must be agreed to and implemented.

Account Balancing: It is up to each participating site to try to maintain a zero balance in the aggregate. Using standard accounting practices, the following scenario is offered as an example. When a site submits a job that will consume GRID resources:

- the resulting resource utilization charges are viewed as a debit to the submitting (consumer) site. The consumer's home site can then decide how to charge the users authorized project and individual account.
- the resulting resource utilization charge is handled as a credit to the resource provider (supplier) site. This entitles jobs at the supplier's site to use an equivalent amount of GRID resources at the consumers site.

Ultimately, the credits provide at a supplier site should be balanced by debits incurred as usage at other GRID sites. A resource supplier could potentially increase demand for its resources (and gather more GRID credits) by lowering its rates and reduce GRID demand by raising its resource rates.

GRID Accounting Policies – Making the Market Work

To create a single metric that represents resource usage, a formula is required that will map from a set of resource utilization or allocation figures to a fixed resource unit. This resource unit will then comprise the basic GRID allocation unit. Project Administrators will be provided with a number of resource allocation units that they can distribute between resources that are made available to them from a GRID producer. The weighting factors in the formula for each GRID site should reflect the supply/demand curve for that resource at that computing site. For example, if a GRID producer provides a high-speed solid-state storage-device resource that is in great demand at their site, they should have a large weighting factor for the use of that resource in the cost formula. Correspondingly, if a GRID producer has a resource that is underused, the weighting factor should be very small for that resource. The utilization of supply and demand curves for all GRID resource will help to create a resource market force that will equalize and load-balance the use of resource across the entire GRID.

The establishment of policies and procedures to handle GRID sites that will not or can not maintain an

aggregate zero balance must be established prior to the production implementation of this accounting system. If a site doesn't have resources to offer the GRID and just want to consume GRID resources, they could establish a partnership with an existing or new GRID member that can "sell" enough resources to balance their demands. In the case where a contributing site can not maintain an aggregate zero balance, a method must be established to transfer balances from one site to another, possibly involving the transfer of real funds. It is also possible for a group of sites to establish a GRID partnership block to declare a single site for balancing purposes.

When examining the question of accounting balance between GRID computing sites, there are several factors to consider. The first factor is what the GRID "fiscal year" or fiscal period ought to be. If it is too small, unnecessary time and effort is spent attempting to "close the books" to achieve some parity between sites heavy with GRID consumers and sites that heavily supply GRID cycles. If the period is too long, large disparities will arise between GRID consumers and producers, and may jeopardize the willingness of GRID producers to participate in a computational GRID. Thus, a suitable "GRID fiscal year" must be identified that will allow the GRID participants to close the books and address any disparities between GRID consumers and producers. How the disparities will be taken care of is not yet defined – it may require the transfer of real funds between net consumers and producers, or it may result in the reduction of resources made available to a group of consumers.

Required Functionality to Provide Services

The basic functionality required to provide distributed accounting is the ability to move the accounting information, resource allocation, resource quotes, and resource account management information from the resource provider to the representative of the resource consumer or user. As mentioned above, if it is assumed that the requesting site has an allocation for the user the job is being submitted from, then the requesting site must have a mechanism to obtain the maximum cost of a job prior to submission to the GRID to verify adequate resources. The resource requirement information is generated at the client (or resource consumer) site, then there must be a mechanism for gathering and accepting resource quotes prior to job submission. There must also be a mechanism to provide the accounting information to the requesting site.

References

- [1] Foster, I., and Kesselman, C. (editors), The Grid: Blueprint for a New Computing Infrastructure, Morgan Kaufmann Publishers, 1999.